



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/718,401	11/21/2003	Chatschik Bisdikian	YOR920030233US1	9742
35060	7590	09/14/2007	EXAMINER	
THE LAW OFFICE OF IDO TUCHMAN 82-70 BEVERLY ROAD KEW GARDENS, NY 11415			ZHE, MENG YAO	
ART UNIT		PAPER NUMBER		
		2195		
MAIL DATE		DELIVERY MODE		
09/14/2007		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

}

Office Action Summary	Application No.	Applicant(s)
	10/718,401	BISDIKIAN ET AL.
	Examiner	Art Unit
	MengYao Zhe	2109

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 21 November 2003.
 2a) This action is **FINAL**. 2b) This action is non-final.
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-35 is/are pending in the application.
 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
 5) Claim(s) _____ is/are allowed.
 6) Claim(s) 1-35 is/are rejected.
 7) Claim(s) _____ is/are objected to.
 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.
 10) The drawing(s) filed on 21 November 2003 is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO/SB/08)
 Paper No(s)/Mail Date _____

4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____
 5) Notice of Informal Patent Application
 6) Other: _____

DETAILED ACTION

1. Claims 1-35 are presented for examination.

Claim Objections

2. Claim 25 objected to because of the following informalities:
 - i) line 3, instead of codes “coupled” to the tangible media, the applicant should change it to codes “stored” to the tangible media. Appropriate correction is required.

Claim Rejections - 35 USC § 112

3. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

4. Claims 1-35 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

- A. The following terms lack proper antecedent basis:

- i) Claim 1, lines 1, “the work load” lack antecedent basis.

Claims 15 and 25 have the same deficiencies as claim 1 above.

- ii) Claim 13, line 4, “the workload elements” lack antecedent basis.

iii) Claim 14, line 22, "the load-dependent group" lack antecedent basis.

B. The claim languages in the following claims are unclear and indefinite:

i) Claim 1, it is unclear what the relationship is between "the workload of an active resource" in lines 1-2 and "parent workload group" < i.e. is the parent workload group part of the workload of an active resource?>

it is unclear what the relationships are among "the workload of an active resource" in lines 1-2, "a collection of units" in line 2, and "a collection of workload units" < i.e. while it is clear that "a collection of units" belong to "the workload of an active resource", it is unclear if "a collection of workload units" are also part of the "workload of an active resource". Is "a collection of units" the same as "a collection of workload units"?>

lines 6-7, it is unclear what is meant by "the parent workload group share an identical sequence of values are a specified depth value of their key identifiers" < i.e. how is the depth specified? How is the depth measured and determined? What is considered to be the depth of key identifiers?>

lines 8-9, it is unclear what is meant by "the identical sequence of values defining a group key identifier associated with the parent workload group" < i.e. does the parent workload group have its own key identifier just like its workload units or it only has a group key identifier as its key

identifier? What is the difference between a group key identifier and a key identifier? How does the identical sequence of value actually define a group key identifier? Is the identical sequence of value exactly the same as the group key identifier?>

lines 13-14, it is unclear what is meant by "increasing the depth value of the parent workload group..." < i.e. Since depth values are used to describe key identifiers as claimed in lines 6-7, while it is reasonable to increase the depth value of key identifiers, it is unclear how the parent workload group itself be increased in depth value.>. Furthermore, it is unclear as to what is meant by "two child workload groups are identified" < i.e. are two child workload groups dynamically created?>. It is unclear what the relationship is between the "two child workload groups" in line 14 and the "workload units" of the parent workload in line 5 <i.e. if the parent is already defined by "workload units" exclusively, how can two more child workload groups be created? Are the child workload groups a further partition of one of the workload units that belongs to the parent?>. It is also uncertain what the relationship is between "the parent workload" in line 13 and "at least one parent workload group" in line 4 < i.e. in the case that there are two parent workload groups associated with an active resource, how does one decide which parent workload group out of the two groups gets its depths value increased?>

lines 15-16, it is unclear how a target resource gets assigned <i.e. is it a target resource that does not have an overload condition that gets assigned?>. It is unclear how the "target resource" is related to the "active resource" in line 2 < i.e. are there multiple resources? And if so, is the target resource a different resource from the active resource? Furthermore, does the target resource have its own associated workload units, parents, and key identifiers?>. It is uncertain as to what is meant by "manage at least one of the child workload groups" <i.e. what is in a workload group? Does it have tasks or jobs? If so, does managing a group mean a group of tasks get to use a set of resources during execution?>

Claims 15 and 25 have the same deficiencies as claim 1 above.

ii) Claim 2, lines 2-3, it is uncertain as to how to identify a "candidate resource" < i.e. is it based on the workload of the candidate resource? The one that is least loaded?>. It is unclear what the relationship is between the "candidate resource" and the "target resource" in line 15, claim 1 <i.e. are they the same?>

Claims 16 and 26 have the same deficiencies as claim 2 above.

iii) Claim 4, it is unclear what is meant by "recording the parent workload group as inactive at the active resource" < i.e. what is meant by an inactive workload group? And what is the purpose of making it inactive if it has an active resource?>

Claims 18 and 28 have the same deficiencies as claim 4 above.

iv) Claim 5, it is unclear how to perform the step of "transferring ..." <i.e. what gets transferred to where? It appears to be that it is the application-specific objects associated with the child workload groups that gets transferred, but where does it transferred? To the target resource mentioned in line 15 of claim 1? Furthermore, do all workload groups or units have application-specific objects associated with them?>

Claims 29 and 19 have the same deficiencies as claim 5 above.

v) Claim 6, it is unclear what portions of "entities operating on elements of the parent workload group" gets redirected < i.e. If parents have workload units that make up the parent, entities operating on the parent must also be operating across the workload units as well. Since claim 1 only claims that at least one child workload group gets redirected, it seems that only portions of entities operating on elements of the redirected child workload groups should get redirected, not the entire entities operating on the entire parent.>

Claim 30 has the same deficiencies as claim 6 above.

vi) Claim 7, lines 2, it is uncertain who is "receiving a probe message" <i.e . is it the active resource that receives this message?>.

lines 4, it is unclear how "guessing a depth..." is done and who is doing the guessing < i.e. a random number is generated? And is it the entity of the workload unit that performs the guessing?>

line 6, it is unclear who is doing the "sending". Furthermore, it is unclear what the relationship is between the "probe message" in line 3 and the "response" in line 6 < i.e. does the probe message help to determine the response?>

Claim 31 has the same deficiencies as claim 7 above.

vii) Claim 8, line 2, it is uncertain how to "refine its estimate" <i.e. how does it know that the estimate is wrong? And how does it make it better or more accurate?>

viii) Claim 9, lines 5-8, it is unclear what the relationship are among "at least two workload groups" in line 5, "consolidated workload group" in line 6, "workload units" in line 7, and "parent workload group" in line 4 of claim 1 < i.e. do the two workload groups belong to the parent workload group? Does the workload units in line 7 belong to the workload groups in line 5? What is the difference between workload units and workload groups? Are they suppose to be the same?>

Claims 20 and 32 have the same deficiencies as claim 9 above.

ix) Claim 10, lines 2-3, it is unclear how "depth value of the parent workload group" may be decreased since in claim 1, line 8, the depth values are used to describe key identifiers instead of the workload group itself.

Claim 21 has the same deficiencies as claim 10 above.

x) Claim 11, line 3, it is unclear what "attributes of the workload unit" may be <i.e. are they the name, load measurement, etc of the workload unit?>

Claims 22 and 34 have the same deficiencies as claim 11 above.

xi) Claim 12, it is unclear how "a virtual key" is constructed. It is unclear what "a load-depended subset of a complete key identifier is < i.e. what is it mean to be a subset of a key identifier? And what does it mean to be load-dependent?>

Claim 23 has the same deficiencies as claim 12 above.

xii) Claim 13, it is unclear how the "mapping service" operates <i.e. what method does it use to return the identity of the target resource?>. It is unclear what the relationship is between the "workload elements" in lines 4 and "workload units" in line 5 of claim 1 < i.e. are they the same thing? How can workload elements belong to a virtual key?>

xiii) Claim 14, lines 7-9, it is unclear how "each individual resource" is related "a set of active resources" < i.e. are the individual resources a sub set of the set of active resources?>.

It is unclear what the relationship is between "a dynamically varying group of key identifiers" in line 9 and "an entire set of key identifiers" in lines 7-8 < i.e. is a group of key IDs in line 9 part of the entire set of key identifiers in lines 7-8?>

Lines 10-12, it is unclear how workloads may be reduced or increased by merely creating or consolidating a group of key identifiers < i.e. workloads generally is a measure of amount of tasks or jobs that are executing, just by changing key identifiers associated with the tasks cannot change the workload alone.>

It is unclear what the relationships are among "an overall set of resources" in line 13, "a set of active resources" in line 7, and "a dynamically varying set of distributed resources" < i.e. are they the same set?>

Lines 16-19, it is uncertain as to what a "virtual key" may be <i.e. what is it? How is it created? And how is it different from a key identifier?>. It is unclear how to "produced the identity of the target resource" merely based on a virtual key < i.e. what is the target resource for and how is it related to the virtual key? And how is its identity produced? Based on a

random selection?>. It is unclear what the relationship is between the "the target resource" and "a set of active resources" defined in line 7 < i.e. is the target resource part of that set?>

It is unclear what the relationship is between the "at least one key identifier" that is associated with a client in line 21 and the "key identifier" that is associated with a workload unit in line 3 <i.e. are they the same kind of identifiers?>

Line 22, it is unclear what the relationship is between "the load-dependent group of identifier keys" and the "a dynamically varying group of key identifiers" in line 9 < i.e. is the dynamically varying group of key identifiers also load dependent? What does it mean for a group of identifier keys to be load-dependent? Are the two groups actually the same group?>

Line 23, it is unclear if "identifier keys" are the same as "key identifier" in line 21. < i.e. if they are the same thing, the same name should be used.>

xiv) Claim 24, it is unclear what the relationship is between the "external service" and the "system" in claim 15 <i.e. what is this external service external to? The system of claim 15? Or is it just external to the resources?>

xv) Claim 33, lines 3, it is unclear how to “decrease the depth value of the parent workload group” while in claim 25, line 9, depth value is used to describe key identifiers, not the workload group itself. Furthermore, it is uncertain as to how “the consolidated workload group is identified” by merely decreasing the “depth value of the parent workload group” <i.e. what is being consolidated? The parent workload group and another workload group? Or the children of the parent workload group? If it is the children of the parent that are being consolidated, the depth of the parent does not decrease.>

xvi) Claim 35, it is uncertain as to what is meant by virtual key includes a hash value <i.e. hash functions or mappings work as follows: a function H takes in input m and outputs value h, $h=H(m)$. h the output, is called the hash value. So what is the input and output in claim 35? Is the virtual key the input, if it is, it can’t include a hash value, since the hash value is the output. What is the output? Is it the key ID or the target resource?>.

Claim Rejections - 35 USC § 102

5. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section

Art Unit: 2109

351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

6. Claims 1, 2, 3-6, 11, 12, 15, 16-19, 22, 23, 25-30, 34 are rejected under 35 U.S.C. 102(e) as being anticipated by Bjornson et al., Pub No. 2002/0194173 (hereafter Bjornson).

7. As per claims 1, 15, and 25, Bjornson teaches the invention as claimed including a method for dynamically adjusting the workload of an active resource, the workload being expressed as a collection of units, each unit including its own key identifier, the active resource being associated with at least one parent workload group, the parent workload group including a collection of workload units such that workload units belonging to the parent workload group share an identical sequence of values at a specified depth value of their key identifiers, the identical sequence of values defining a group key identifier associated with the parent workload group (Fig 5: please note the names assigned to the task and sub-task blocks colored in white), the method comprising:

independently determining by the active resource that an overload condition exists at the active resource (Paragraph 60);

if the overload condition exists:

increasing the depth value of the parent workload group such that at least two child workload groups are identified (paragraphs 60-61);

assigning a target resource to manage at least one of the child workload groups

(paragraphs 56, 60-61).

8. As per claims 2, 16, and 26, Bjornson teaches if the overload condition exists, identifying at least one candidate resource to which the child workload groups may be distributed (Paragraphs 56 and 60: one of the child is kept by the original worker computer, which corresponds to a resource. The other child is added to the VSM bulletin board for another worker computer to take when it is not busy).

9. As per claims 3, 17, and 27, Bjornson teaches requesting workload acceptance from the target resource (Paragraphs 56, 60-61).

10. As per claims 4, 18, and 28, Bjornson teaches recording the parent workload group as inactive at the active resource (Paragraph 56: those worker computer that are forced to wait are considered to be inactive.).

11. As per claims 5, 19, and 29, Bjornson teaches transferring application-specific objects corresponding to the child workload groups (Paragraphs 60-62: tasks are essentially programs that implements algorithms. The Examiner has interpreted the program to be application-specific objects. When one of the divided task is added back to the Task List, the sub-task is transferred.).

12. As per claims 6 and 30, Bjornson teaches redirecting entities operating on elements of the parent workload group to the target resource managing the child workload group (Paragraph 56, 60, 71-74: database corresponds to entities. They can

be split up according to how the tasks are split up. When a sub-task with its associated database gets assigned to another worker computer, entities are redirected.).

13. As per claims 11, 22, and 34, Bjornson teaches further comprising associating the workload unit with the key identifier such that the key identifier encodes one or more attributes of the workload unit (Fig 5: Because the naming system contains the name of the parent task for each sub-task, the parent identification in a child's name corresponds to an attribute.).

14. As per claims 12 and 23, Bjornson teaches constructing a virtual key for mapping to the target resource, wherein the virtual key includes a load-dependent subset of a complete key identifier (Fig 5, unit d: Taskt1.B is the complete key identifies and ".B" is the subset.).

Claim Rejections - 35 USC § 103

15. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

16. Claims 7-10, 13, 14, 20, 21, 24, 31-33, 35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Bjornson et al., Pub No. 2002/0194173 (hereafter Bjornson).

17. As per claims 7 and 31, Bjornson discloses a system that can break down a single task into different levels of sub-tasks, each with its associated identifier, so that they may be assigned to computers according to the computer's dynamic workload. Furthermore, Bjornson teaches estimating the amount of computational resources available in each computer to see if a task needs to be broken down. Bjornson also teaches a group key identifier that indicates the nearest known active parent group to which it belongs (Fig 5; Paragraphs 60-61).

18. Bjornson, however is silent as to, the specifics of receiving a probe message from an entity operating on a workload unit that is a member of the parent workload group, the probe message including a guessed identifier key formed by guessing a depth to be associated with the unit's key identifier; and sending a response to the entity indicating the group key identifier that the current resource locally determines to be the nearest known active parent group to which the element's key identifier belongs.

19. To specifically specify the detailed steps of receiving a probe message from an entity operating on a workload unit that is a member of the parent workload group, the probe message including a guessed identifier key formed by guessing a depth to be associated with the unit's key identifier; and sending a response to the entity indicating the group key identifier that the current resource locally determines to be the nearest known active parent group to which the element's key identifier belongs would have been obvious to one of ordinary skill in the art as it is well known that in order for the worker computer to keep track of the tasks and subtasks it is working on, it needs to keep track of the parent groups of all of its child task groups.

20. As per claim 8, Bjornson teaches wherein the entity operating on a workload unit uses the response to further refine its estimate of a correct depth to be associated with the unit's key identifier; and probing another resource associated with the parent key group formed by using the refined depth of the unit's key identifier (Paragraphs 60-61).

21. As per claims 9, 20, and 32, Bjornson teaches

determining that an under-load condition exists at the active resource (Paragraphs 95-96: limitation is set for the amount of time a worker computer is to remain idle, when it is exceeded, it is inherent that the system will know that a worker computer is under-loaded.)

identifying at least two workload groups for consolidation into a consolidated workload group; generating a consolidated key identifier such that workload units belonging to the consolidated workload group share an identical sequence of values at a specified depth value of the consolidated key identifier; and managing the consolidated workload group by the active resource (Paragraph 69 and Fig 5).

22. Bjornson does not teach consolidation of workload group specifically under the condition that there is an under-load at the active resource. However, it would have been obvious to one having ordinary skill in the art at the time of the applicant's invention to combine the teachings of consolidating workload groups and determining an under-load condition exists so that under the specific situation that an under-load condition exists, workload groups are consolidated, because just like Bjornson's

teaching of increasing the depth value of parent workload group to balance the workload in the situation that one workload group is overloaded, consolidation of workload group is another way to balance the workload group in the opposite situation.

23. As per claims 10 and 21, Bjornson teaches wherein generating the consolidated key identifier includes decreasing the depth value of the parent workload group such that the consolidated workload group is identified (Fig 5: please note the ID given for task blocks colored in gray.).

24. As per claim 13, Bjornson does not specifically teach using the constructed load-dependent virtual key as an input to a separate mapping service that returns the identity of the target resource to which the workload elements belonging to the virtual key should be directed.

25. However, it would have been obvious to one having ordinary skill in the art at the time of the invention to let each worker computer, which corresponds to the resources, keep track of the tasks that are associated with its task ID so that it knows what part of an overall task it is currently working on, and further so that it may reconsolidate the results of each sub-tasks into one final task result in the end (Fig 5, paragraphs 63 and 69.).

26. As per claim 14, Bjornson teaches A system for running a distributed computer application whose workload can be decomposed into a set of workload units, each workload unit including its own key identifier, over a dynamically varying set of distributed resources, the number of resources involved in the distributed computation

varying dynamically in response to changes in an overall workload, the system comprising (Fig 5 and paragraph 60):

a set of active resources cooperatively managing an entire set of key identifiers constituting the overall workload, each individual resource managing a dynamically varying group of key identifiers, each resource independently evaluating its own workload condition and deciding on the creation to reduce its workload (Fig 5, paragraphs 60);

an overall set of resources, of which the active resources constitute a subset that can be utilized as part of the distributed computer application as needed (Paragraphs 60-61);

a set of client entities utilizing the distributed application, each client entity being associated with at least one key identifier, and each client entity dynamically determining the load-dependent group of identifier keys that each such currently belongs to (Fig 5).

27. Bjornson does not teach consolidation of group key identifiers to increase its workload and a mapping service configured to receive a virtual key associated with at least one key identifier as input and configured to produce the identity of the target resource from the overall resource set as an output.

28. However, it would have been obvious to one having ordinary skill in the art at the time of the applicant's invention to combine the teachings of consolidating workload groups and determining an under-load condition exists so that under the specific

situation that an under-load condition exists, workload groups are consolidated, because just like Bjornson's teaching of increasing the depth value of parent workload group to balance the workload in the situation that one workload group is overloaded, consolidation of workload group is another way to balance the workload group in the opposite situation. It is also obvious that each worker computer, which corresponds to the resources, keep track of the tasks that are associated with its task ID so that it knows what part of an overall task it is currently working on, and further so that it may reconsolidate the results of each sub-tasks into one final task result in the end (Fig 5, paragraphs 63 and 69.).

29. As per claim 24, Bjornson teaches an external service configured to identify at least one candidate resource to which the child workload groups may be distributed (Paragraph 56: the VSM, which are external to the worker computers governs how the task may be taken by the worker.).

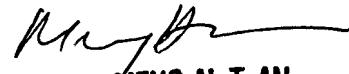
30. As per claim 33, Bjornson teaches wherein the program code to generate the consolidated key includes program code to decrease the depth value of the parent workload group such that the consolidated workload group is identified (Fig 5 and Paragraph 69).

31. As per claim 35, Bjornson teaches program code configured to construct a virtual key for mapping to the target resource, wherein the virtual key includes a hash value of the key identifier (Fig 5).

Conclusion

32. Any inquiry concerning this communication or earlier communications from the examiner should be directed to MengYao Zhe whose telephone number is 571-272-6946. The examiner can normally be reached on Monday Through Friday, 7:30 - 5:00 EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Meng-Ai An can be reached at 571-272-3756. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.



MENG-AL T. AN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100